

AMENDMENTS TO THE CLAIMS

This listing of claims will replace all prior versions, and listings, of claims in the application:

Listing of Claims:

- 1 1. (Currently Amended) A method of flow controlling InfiniBand receive traffic, comprising:
 - 3 maintaining a single memory structure for queuing InfiniBand traffic
 - 4 received via multiple virtual lanes and multiple queue pairs;
 - 5 identifying a first packet payload received via a first virtual lane and a first queue pair;
 - 7 determining whether the first payload can be stored in the memory structure without exceeding a portion of the memory structure allocated to the first virtual lane;
 - 10 determining whether the first payload can be stored in the memory structure without exceeding a portion of the memory structure allocated to the first queue pair; ~~and~~
 - 13 if storing the first payload in the memory structure would exceed said portion of the memory structure allocated to the first queue pair, determining whether the first queue pair is enabled to use a shared portion of the memory structure to store payloads of packets received via the first queue pair; and
 - 17 maintaining a second memory configured to store, for each of the multiple queue pairs that is active, one or more parameters associated with operation of said queue pair, wherein said parameters include:
 - 20 a maximum number of message credits advertisable by said queue pair;

a maximum number of memory structure buffers dedicated to storing payloads of packets received via said queue pair;
an indicator configured to indicate whether said queue pair is enabled to use a set of shared memory structure buffers; and
a number of shared memory structure buffers in said set of shared memory structure buffers, wherein said shared memory structure buffers are available for use by said queue pair to store payloads of packets received via said queue pair if:
said queue pair has used said maximum number of memory structure buffers; and
said indicator indicates that said queue pair is enabled to use said set of shared memory structure buffers; and
a maximum number of message credits advertisable by said queue pair when said queue pair starts using said shared memory structure buffers.

2. (Original) The method of claim 1, further comprising:
 2. allocating a portion of the memory structure to each of the multiple virtual lanes; and
 4. allocating a portion of the memory structure to each of the multiple queue pairs.

3. (Original) The method of claim 1, wherein the memory structure comprises a set of linked lists of memory structure buffers, including one linked list for each of the multiple queue pairs that are active.

1 4. (Original) The method of claim 1, further comprising:
2 dropping the first payload if the first payload cannot be stored in the

3 memory structure without exceeding the portion of the memory structure allocated
4 to the first virtual lane.

1 5. (Original) The method of claim 1, further comprising:
2 issuing a Retry, Not Ready, Negative Acknowledgement (RNR-NAK) if:
3 the first payload cannot be stored in the memory structure without
4 exceeding a portion of the memory structure allocated to the first queue
5 pair; and
6 the first queue pair is not enabled to use the shared portion of the
7 memory structure.

1 6. (Original) The method of claim 1, further comprising:
2 issuing a Retry, Not Ready, Negative Acknowledgement (RNR-NAK) if:
3 the first payload cannot be stored in the memory structure without
4 exceeding a portion of the memory structure allocated to the first queue
5 pair;
6 the first queue pair is enabled to use the shared portion of the
7 memory structure; and
8 the shared portion of the memory structure is full.

1 7. (Original) The method of claim 1, further comprising:
2 defining one or more dedicated thresholds in the portion of the memory
3 structure allocated to the first queue pair; and
4 for each of said dedicated thresholds, identifying a number of message
5 credits the queue pair may advertise when the amount of the memory structure
6 used by the queue pair exceeds said dedicated threshold.

1 8. (Original) The method of claim 1, further comprising:

2 defining one or more shared thresholds in the shared portion of the
3 memory structure; and
4 for each of said shared thresholds, identifying a number of message credits
5 the queue pair may advertise when the amount of the shared portion used by the
6 multiple queue pairs exceeds said shared threshold.

1 9. (Original) The method of claim 1, further comprising:
2 receiving a request on a second queue pair to perform an RDMA (Remote
3 Direct Memory Access) Read operation; and
4 based on an amount of data expected to be received via the RDMA Read
5 operation, reserving a sufficient number of buffers in the memory structure.

1 10. (Original) The method of claim 1, further comprising:
2 in the single memory structure, reassembling the queued InfiniBand traffic
3 into outbound communications;
4 receiving a payload on an idle queue pair, wherein a queue pair is idle if no
5 traffic from the queue pair is stored in the single memory structure; and
6 only queuing the payload in the single memory structure if sufficient space
7 in the single memory structure is reserved for completing reassembly of outbound
8 communications on each non-idle queue pair.

1 11. (Currently Amended) A computer readable medium storing
2 instructions that, when executed by a computer, cause the computer to perform a
3 method of flow controlling InfiniBand receive traffic, the method comprising:
4 maintaining a single memory structure for queuing InfiniBand traffic
5 received via multiple virtual lanes and multiple queue pairs;
6 identifying a first packet payload received via a first virtual lane and a first
7 queue pair;

8 determining whether the first payload can be stored in the memory
9 structure without exceeding a portion of the memory structure allocated to the first
10 virtual lane;

11 determining whether the first payload can be stored in the memory
12 structure without exceeding a portion of the memory structure allocated to the first
13 queue pair; and

14 if storing the first payload in the memory structure would exceed said
15 portion of the memory structure allocated to the first queue pair, determining
16 whether the first queue pair is enabled to use a shared portion of the memory
17 structure to store payloads of packets received via the first queue pair; and
18 maintaining a second memory configured to store, for each of the multiple
19 queue pairs that is active, one or more parameters associated with operation of
20 said queue pair, wherein said parameters include:

21 a maximum number of message credits advertisable by said queue
22 pair;

23 a maximum number of memory structure buffers dedicated to
24 storing payloads of packets received via said queue pair;

25 an indicator configured to indicate whether said queue pair is
26 enabled to use a set of shared memory structure buffers; and

27 a number of shared memory structure buffers in said set of shared
28 memory structure buffers, wherein said shared memory structure buffers
29 are available for use by said queue pair to store payloads of packets
30 received via said queue pair if:

31 said queue pair has used said maximum number of memory
32 structure buffers; and

33 said indicator indicates that said queue pair is enabled to
34 use said set of shared memory structure buffers; and
35 a maximum number of message credits advertisable by said queue

36 pair when said queue pair starts using said shared memory structure
37 buffers.

1 12. (Original) The computer readable medium of claim 11, wherein the
2 method further comprises:

3 defining one or more dedicated thresholds in the portion of the memory
4 structure allocated to the first queue pair; and
5 for each of said dedicated thresholds, identifying a number of message
6 credits the queue pair may advertise when the amount of the memory structure
7 used by the queue pair exceeds said dedicated threshold.

1 13. (Original) The computer readable medium of claim 11, wherein the
2 method further comprises:

3 defining one or more shared thresholds in the shared portion of the
4 memory structure; and
5 for each of said shared thresholds, identifying a number of message credits
6 the queue pair may advertise when the amount of the shared portion used by the
7 multiple queue pairs exceeds said shared threshold.

1 14. (Original) The computer readable medium of claim 11, wherein the
2 method further comprises issuing a Retry, Not Ready, Negative
3 Acknowledgement (RNR-NAK) only if one of:

4 (a) the first payload cannot be stored in the memory structure without
5 exceeding a portion of the memory structure allocated to the first queue
6 pair; and
7 the first queue pair is not enabled to use the shared portion of the
8 memory structure; and
9 (b) the first payload cannot be stored in the memory structure without

10 exceeding a portion of the memory structure allocated to the first queue
11 pair;
12 the first queue pair is enabled to use the shared portion of the
13 memory structure; and
14 the shared portion of the memory structure is full.

1 15-29. Cancelled

1 30. (Currently Amended) A method of avoiding locking, in receive
2 InfiniBand queues, the method comprising:
3 maintaining a single memory structure for reassembling InfiniBand traffic
4 received via multiple virtual lanes and multiple queue pairs;
5 identifying a first packet payload received via a first queue pair that is idle,
6 wherein the first queue pair is considered idle if no traffic from the first queue pair
7 is stored in said single memory structure;
8 for each other queue pair for which traffic from said queue pair is stored in
9 said single memory structure, determining whether sufficient space in the single
10 memory structure is reserved for reassembling said traffic; and
11 storing the first packet payload in said single memory structure only if
12 sufficient space in the single memory structure is available for reassembling said
13 traffic; and
14 maintaining a second memory configured to store, for each of the multiple
15 queue pairs that is active, one or more parameters associated with operation of
16 said queue pair, wherein said parameters include:
17 a maximum number of message credits advertisable by said queue
18 pair;
19 a maximum number of memory structure buffers dedicated to
20 storing payloads of packets received via said queue pair;

21 an indicator configured to indicate whether said queue pair is
22 enabled to use a set of shared memory structure buffers; and
23 a number of shared memory structure buffers in said set of shared
24 memory structure buffers, wherein said shared memory structure buffers
25 are available for use by said queue pair to store payloads of packets
26 received via said queue pair if:
27 said queue pair has used said maximum number of memory
28 structure buffers; and
29 said indicator indicates that said queue pair is enabled to
30 use said set of shared memory structure buffers; and
31 a maximum number of message credits advertisable by said queue
32 pair when said queue pair starts using said shared memory structure
33 buffers.

1 31. (Original) The method of claim 30, wherein said determining
2 comprises, for each said other queue pair:
3 identifying an amount of space in said single memory structure reserved
4 for said other queue pair; and
5 comparing said amount of reserved space to an amount of space expected
6 to be needed to complete reassembly of said traffic from said other queue pair.

32. (Currently Amended) An apparatus for flow controlling received
InfiniBand traffic, comprising:
a single memory structure configured to queue payloads of InfiniBand
traffic received via multiple virtual lanes and multiple queue pairs;
a resource manager configured to manage the memory structure;
a first module configured to facilitate the advertisement of virtual lane
credits; and

8 a second module configured to facilitate the advertisement of queue pair
9 credits; and credits;

10 a second memory configured to store, for each of the multiple queue pairs
11 that is active, one or more parameters associated with operation of said queue
12 pair, wherein said parameters include:

13 a maximum number of message credits advertisable by said queue
14 pair;

15 a maximum number of memory structure buffers dedicated to
16 storing payloads of packets received via said queue pair;

17 an indicator configured to indicate whether said queue pair is
18 enabled to use a set of shared memory structure buffers; and

19 a number of shared memory structure buffers in said set of shared
20 memory structure buffers, wherein said shared memory structure buffers
21 are available for use by said queue pair to store payloads of packets
22 received via said queue pair if:

23 said queue pair has used said maximum number of memory
24 structure buffers; and

25 said indicator indicates that said queue pair is enabled to
26 use said set of shared memory structure buffers; and
27 a maximum number of message credits advertisable by said queue
28 pair when said queue pair starts using said shared memory structure
29 buffers.

1 33. (Original) The apparatus of claim 32, wherein said single memory
2 structure comprises multiple linked lists of memory structure buffers, including
3 one linked list for each of the multiple queue pairs that is active.

1 34. (Original) The apparatus of claim 32, wherein said first module

2 comprises an InfiniBand link core.

1 35. (Original) The apparatus of claim 32, wherein said second module
2 comprises an acknowledgement generator configured to generate transport layer
3 acknowledgements.

1 36. (Original) The apparatus of claim 32, further comprising a
2 processor interface configured to facilitate the programming of operating
3 parameters associated with the multiple virtual lanes and the multiple queue pairs.

1 37. (Original) The apparatus of claim 32, further comprising:
2 a first memory configured to store one or more parameters associated with
3 operation of a first virtual lane.

1 38. (Original) The apparatus of claim 37, wherein said one or more
2 parameters include:

3 a count of the number of memory structure buffers currently used to store
4 payloads of packets received via the first virtual lane; and
5 a threshold, wherein a first packet is dropped if storing the payload of the
6 first packet would cause said count to exceed said threshold.

1 39. (Cancelled)

1 40. (Cancelled)

1 41. (Original) The apparatus of claim 40, wherein said one or more
2 parameters further include:
3 one or more dedicated thresholds, wherein each said dedicated threshold

4 identifies a subset of said maximum number of memory structure buffers; and
5 for each said dedicated threshold, a number of message credits
6 advertisable by said queue pair when said queue pair uses said subset of said
7 maximum number of memory structure buffers.

1 42. (Cancelled)

1 43. (Currently Amended) The apparatus of claim 32, claim 42, wherein
2 said one or more parameters further include:
3 one or more shared thresholds, wherein each said shared threshold
4 identifies a subset of said number of shared memory structure buffers; and
5 for each said shared threshold, a number of message credits advertisable
6 by said queue pair when said queue pair uses said subset of said number of shared
7 memory structure buffers.

1 44-55. (Cancelled)